Artificial Intelligence: Unveiling Hidden Patterns in Human Behaviour

Douglas C. Youvan

doug@youvan.com

July 16, 2023

As artificial intelligence (AI) continues to evolve, its potential for understanding complex human behaviours expands. Beyond simply processing vast amounts of data, AI can unearth non-obvious patterns, revealing fresh insights into human psychology and sociology. This essay explores AI's capabilities in detecting intricate behavioural patterns, its applications in social sciences, the associated challenges and ethical considerations, and predictions for the future. It serves as a comprehensive examination of the transformative role AI could play in deepening our understanding of the intricate tapestry of human behaviour.

Keywords: Artificial Intelligence, Human Behaviour, Social Sciences, Pattern Detection, Data Analysis, Case Studies, Ethical Considerations, Data Privacy, Misuse of AI, Future Trends, Big Data, AI Interpretability, Societal Impact, AI Determinism, AI and Ethics.

Introduction

In the burgeoning field of artificial intelligence, the ability to uncover patterns invisible to the human mind offers an exciting frontier of exploration. With an ever-increasing influx of data in our digital age, AI's potential to process and analyze these vast datasets can provide insights into complex phenomena that would otherwise remain elusive. One area where AI's pattern detection capabilities could be particularly transformative is in understanding human behaviour—a realm that has long fascinated and perplexed psychologists, sociologists, and anthropologists alike.

Human behaviour, characterized by its intricacy and shaped by an array of socio-cultural, economic, and psychological factors, has traditionally been studied through a variety of methodologies, ranging from observational studies to complex statistical analyses. Yet, despite significant strides in these fields, the unpredictability and dynamism of human behaviour often defy simple categorization or comprehension. The latent patterns in human behaviour are often so complex and multi-dimensional that they escape conventional analytical approaches.

Herein lies the significance of our study. As we venture into the labyrinth of human behaviour, AI has the potential to serve as our guide, decoding patterns beyond the reach of human cognition and traditional methodologies. Whether it's predicting social trends, understanding consumer behaviour, or exploring the underpinnings of political sentiment, AI can offer a new lens through which we view the rich tapestry of human experience.

The application of AI in understanding human behaviour is not without its challenges, and it raises compelling questions around feasibility, interpretability, and ethics. Nonetheless, the potential impact is enormous. By shedding light on these hidden patterns,

we can foster a more nuanced understanding of our societies, inform policies, drive innovation, and potentially, address some of the pressing challenges of our times. This exploration, therefore, is not merely a theoretical endeavour; it represents a step towards unlocking untapped potential and creating a more comprehensible world.

Artificial Intelligence: A Brief Overview

Artificial Intelligence (AI), at its core, refers to the simulation of human intelligence in machines programmed to think like humans and mimic their actions. The term also applies to any machine that exhibits traits associated with the human mind, such as learning and problem-solving. AI's power resides in its capacity to analyze, understand, and learn from vast amounts of data, detect patterns, make predictions, and adapt to new information.

AI is a broad field with many sub-areas, but at a high level, it can be categorized into two main types: narrow AI, which is designed to perform a specific task, such as voice recognition, and general AI, which can understand, learn, and apply knowledge across a wide array of tasks, much like a human would. As of my knowledge cutoff in September 2021, narrow AI is a reality, while general AI remains largely theoretical.

The concept of AI traces its roots back to antiquity, with myths, stories, and speculations about artificial beings endowed with intelligence or consciousness by master craftsmen. However, as a scientific discipline, AI is relatively young. It was first proposed as a field of study at the Dartmouth Conference in 1956, where the term 'Artificial Intelligence' was coined. This marked the birth of AI as an autonomous discipline.

The subsequent decades saw a series of ups and downs, with periods of intense excitement and funding—known as AI

summers—followed by periods of disappointment and reduced funding, or AI winters. Early work in AI focused on rule-based systems and symbolic reasoning. However, these approaches were found to be limited when dealing with complex, real-world problems.

The field was revolutionized in the 1980s and 1990s by the introduction of machine learning, where instead of programming explicit rules, computers were taught to learn from data. This revolution culminated in the development of deep learning in the 2000s and 2010s, an approach inspired by the human brain's architecture that uses artificial neural networks to analyze and learn from data. Deep learning has been responsible for many of the most exciting recent advances in AI, from self-driving cars to voice assistants and recommendation algorithms.

Today, AI technologies have infiltrated almost every aspect of our lives and continue to evolve at an unprecedented pace. As we stand on the brink of a new era in AI research and application, it is paramount that we explore and harness its potential to decode the complex patterns in human behaviour, a prospect that forms the crux of our exploration.

The Complexity of Human Behaviour

Human behaviour, a product of millennia of evolution and societal development, is one of the most intricate phenomena on earth. A complex interplay of thoughts, emotions, and actions, it forms the basis of our individual identities and collective societies. To comprehend the magnitude of its complexity, it's crucial to appreciate the multi-faceted nature of human behaviour.

At its core, human behaviour is a reflection of an individual's internal psychological processes—thoughts, feelings, and motivations. These internal processes, which include cognition,

perception, emotion, personality, behaviour, and interpersonal relationships, are inherently complex and nuanced. They can vary greatly among individuals, change over time, and are shaped by a host of biological and environmental influences.

Furthermore, these psychological processes don't operate in a vacuum. They are deeply intertwined with social, cultural, economic, and environmental factors, creating a web of influences that make human behaviour incredibly diverse and dynamic. Each individual's behaviour is a unique confluence of these multiple forces, and these forces themselves are ever-changing, adding to the complexity.

Culture, for instance, significantly shapes human behaviour. It influences our values, beliefs, customs, and even the way we perceive the world. From the food we eat to the clothes we wear, from our rituals to our language, our behaviour is a testament to the cultural milieu in which we are embedded.

Socio-economic status is another crucial determinant. The resources we can access, the education we receive, the neighborhoods we live in, and the opportunities available to us can profoundly influence our behaviours, aspirations, and life trajectories.

Education, both formal and informal, shapes our knowledge, skills, attitudes, and values. It influences how we think, solve problems, communicate, and interact with the world. It can alter our behaviours and perspectives in profound ways.

Geography, too, plays a role in shaping our behaviour. The physical environment, climate, and local resources in the place we live can influence our lifestyles, occupations, diet, and even our attitudes and beliefs.

Biology and personal experiences also have a significant impact. Genetic predispositions, health conditions, early childhood experiences, personal relationships, and individual life events can all leave indelible marks on our behaviour.

When we consider the multitude of factors influencing human behaviour and the ways in which they interact, the complexity is staggering. It's a dynamic, intricate system that is challenging to decipher using traditional methods. However, the advent of AI, with its capability to process and analyze large and complex datasets, opens up new possibilities for understanding this labyrinth of influences and their interplay in shaping human behaviour.

AI in Social Sciences: A New Frontier

The intersection of artificial intelligence and social sciences is a burgeoning field, signaling a paradigm shift in how we understand, analyze, and predict human behaviour. AI offers new tools to social scientists, extending beyond traditional qualitative and quantitative methods, and enabling deeper, more complex analyses.

Currently, AI applications in social sciences are largely centered around data analysis and prediction. Machine learning, a branch of AI, has been particularly impactful in these areas. Machine learning algorithms can process vast amounts of data, learn from them, and make predictions or decisions without being explicitly programmed to perform the task. This characteristic is incredibly useful in areas such as sociology, psychology, economics, and political science, where complex human behaviour is studied.

For instance, AI has been used to analyze social media posts to gauge public sentiment on specific issues, a process known as sentiment analysis. This has applications in areas ranging from

market research, where it's used to understand consumer sentiment, to political science, where it can help gauge public opinion on political issues or candidates.

In psychology, AI has been used to detect patterns in human speech or writing that can indicate mental health conditions. For instance, researchers have used machine learning algorithms to predict depression based on patterns in social media posts.

In economics, AI has been leveraged to predict market trends based on complex interactions of numerous economic indicators. These models can help economists better understand market dynamics and make more accurate forecasts.

While these applications showcase AI's potential in social sciences, it is the prospect of AI unveiling previously unseen patterns in human behaviour that truly represents a new frontier. Traditional methods may not capture the full complexity of human behaviour due to their inability to process and analyze vast amounts of multi-dimensional data. However, AI, with its advanced computational power and sophisticated algorithms, can handle this complexity.

AI can analyze huge datasets from various sources such as social media, census data, economic indicators, and more, in a way that's impossible for human researchers. It can recognize patterns across these datasets, identifying correlations and causations that might elude traditional analyses. These patterns can provide novel insights into the dynamics of human behaviour, enhancing our understanding and prediction of social phenomena.

However, realizing this potential of AI in social sciences is not without challenges. Concerns around data privacy, ethical use of AI, and the interpretability of AI models must be addressed. Moreover, it's essential to remember that AI is a tool that complements, rather than replaces, traditional social science

methods. It is through the synergy of AI and traditional methodologies that we can hope to unravel the complexities of human behaviour in its full depth and richness.

Limitations of Traditional Approaches

In the realm of social sciences, traditional methods used to analyze human behaviour have primarily fallen into two categories: qualitative and quantitative.

Qualitative research methodologies, such as interviews, focus groups, ethnographies, and case studies, aim to gather an in-depth understanding of human behaviour and the reasons that govern such behaviour. They provide rich, detailed data that give us insights into individuals' thoughts, feelings, and experiences. These methods excel in their capacity to provide context and capture the nuance of human experience.

Quantitative research methods, on the other hand, involve systematic empirical investigation through statistical, mathematical, or computational techniques. Surveys, experiments, and observational studies fall into this category. These methods allow researchers to measure and analyze data numerically, making them powerful tools for testing hypotheses and drawing generalizable conclusions.

Despite their considerable merits, traditional methods come with certain limitations, especially when it comes to detecting complex patterns in human behaviour.

For qualitative methods, the subjective nature of data collection and interpretation can lead to bias. While these methods provide depth, they often lack breadth due to the typically small sample sizes. Additionally, the complexity and richness of qualitative data

make it difficult to identify underlying patterns, particularly when those patterns involve interactions between numerous variables.

Quantitative methods, while effective in handling large datasets and providing generalizable conclusions, often operate under strict assumptions about the nature of data and relationships between variables. Real-world data often violate these assumptions, making the models less accurate. Moreover, conventional statistical methods may struggle to capture nonlinear relationships or interactions between a large number of variables, which are common in human behaviour data.

Another significant limitation of traditional methods is their inability to effectively process and analyze the vast amounts of data generated in our digital age. Social media posts, online transactions, GPS location data, health records, among others, represent a goldmine of information about human behaviour. Traditional methods are ill-equipped to handle such Big Data, with its volume, variety, and velocity.

Additionally, traditional methods often analyze data in a reductionist manner, simplifying complex phenomena into manageable bits. While this approach makes the analysis tractable, it risks losing sight of the complexity of human behaviour, where many variables interact simultaneously in intricate ways.

In contrast, artificial intelligence, with its ability to handle Big Data, detect complex patterns, and learn from data without explicit programming or rigid assumptions, can potentially overcome these limitations. This makes AI a valuable tool for social scientists aiming to understand the labyrinthine world of human behaviour. However, this potential should be harnessed carefully, taking into account the ethical and practical challenges that the use of AI brings along.

Harnessing AI to Unveil Hidden Patterns

Artificial intelligence, specifically machine learning and deep learning, presents unprecedented opportunities to discover non-obvious patterns in complex datasets. AI can identify multi-dimensional correlations, interactions between variables, and non-linear relationships that might go unnoticed by traditional analysis methods.

Machine learning algorithms excel at finding patterns in data. By exposing these algorithms to a variety of inputs and outcomes, they learn to make connections and predictions. Deep learning, a subset of machine learning, takes this even further by using artificial neural networks inspired by the human brain. These networks can process and learn from multi-layered, high-dimensional data, finding patterns too complex or subtle for humans or simpler algorithms to detect.

The ability of AI to handle big data enables it to tap into an extensive array of information sources, ranging from social media posts, text messages, online search data, and digital transactions to physiological data from wearable devices. This capability allows AI to analyze data across multiple dimensions, resulting in a holistic understanding of human behaviour that wouldn't be possible with a single data source or traditional analytical techniques.

Consider a hypothetical example in the realm of public health. Suppose there's a growing public health concern, such as an increase in mental health issues among the youth. Traditional approaches might examine this problem by looking at individual factors, such as socio-economic status, school performance, or family history. While these analyses are valuable, they may not capture the entire picture.

Here, AI could contribute by analyzing a broad array of data, including social media posts, browsing history, mobile phone usage, physical activity, school records, and more. By analyzing this data collectively, AI could identify hidden patterns and interactions between these variables that contribute to mental health issues.

As another example, consider the understanding of political behaviour. AI could analyze vast amounts of data, including voter demographics, social media discourse, economic indicators, local news coverage, and more, to predict election outcomes or public support for policies. It might reveal complex interplays between these factors that traditional methods could overlook.

In the world of marketing, AI could analyze consumer behaviour by considering not just obvious factors like past purchases and demographic characteristics, but also social media activity, product reviews, online search trends, and more. This could lead to the discovery of complex consumer segments and purchase triggers, aiding the development of highly targeted marketing strategies.

These hypothetical examples illustrate how AI could unearth complex, non-obvious patterns in human behaviour. It's important to note, though, that while AI has immense potential, its implementation needs to be done carefully, respecting privacy rights and ensuring interpretability of the results. By judiciously leveraging AI, we can significantly augment our understanding of human behaviour, enabling informed decision-making in various societal domains.

Case Studies: AI in Action

The use of artificial intelligence to discern intricate patterns in human behaviour is no longer a futuristic idea; it's actively

shaping various sectors today. The following real-world case studies illuminate how AI has been employed to identify complex patterns, the insights that emerged from these endeavours, and their far-reaching implications.

Case Study 1: Consumer Behaviour

AI has been instrumental in revolutionizing the e-commerce sector, particularly in understanding and predicting consumer behaviour. A notable example is Amazon, the world's largest online marketplace. Amazon utilizes AI-powered recommendation systems to suggest products to users based on their browsing history, past purchases, and products viewed or purchased by similar customers. This system leverages complex patterns of consumer behaviour, considering factors far more numerous and subtle than a human could manually account for.

The impact of these insights is substantial, not only driving increased sales for Amazon but also improving the shopping experience for consumers by personalizing it to their preferences.

Case Study 2: Political Sentiment Analysis

AI has been used to analyze and predict political sentiment, notably during election seasons. A notable case was during the 2016 U.S. Presidential Elections, where an AI system called MogIA correctly predicted the outcome based on 20 million data points from platforms like Google, YouTube, and Twitter. MogIA used sentiment analysis to identify patterns in the public's opinion about the candidates, helping to predict voting behaviour more accurately than many traditional polls.

This example underscores AI's potential to predict political sentiment and highlights the importance of analyzing diverse data sources to capture the complex dynamics of public opinion.

Case Study 3: Mental Health Monitoring

AI has made strides in mental health research, particularly in early detection of conditions such as depression or anxiety. An example of this is Facebook's AI-powered tool that scans posts and comments for signs of suicidal ideation. By detecting patterns in users' writings that may indicate distress, the system alerts a team of human reviewers who can intervene if necessary.

The tool has the potential to provide timely help to individuals who might not have sought it themselves, illustrating how AI can serve as a valuable tool in mental health intervention and prevention.

Case Study 4: Traffic Management

In urban planning and management, AI has been employed to optimize traffic flow in cities. For instance, Google's AI company, DeepMind, partnered with Google Maps to predict traffic and suggest the fastest routes to users. The system processes vast amounts of real-time data, including traffic conditions, construction zones, accidents, and even weather conditions, to identify patterns and make accurate predictions.

This application of AI helps reduce travel time, fuel consumption, and emissions, showcasing AI's potential in addressing environmental and urban planning challenges.

These case studies demonstrate the significant potential of AI in unveiling complex patterns and contributing to diverse domains, from e-commerce and politics to mental health and urban planning. However, they also underscore the need for responsible use of AI, considering privacy, consent, and the societal implications of these technologies.

Challenges in Utilising AI for Pattern Detection

While artificial intelligence presents significant opportunities for unearthing intricate patterns in human behaviour, it also poses unique challenges. From issues related to the interpretability of AI models to concerns about data privacy, the implementation of AI for pattern detection in social sciences needs careful consideration and management.

One of the principal challenges in using AI, particularly machine learning and deep learning models, is the 'black box' problem. This term refers to the opacity of these models; while they are capable of making incredibly accurate predictions or discerning complex patterns, the internal workings of these models are often not easily interpretable. This opacity can make it difficult to understand how the AI system arrived at a particular conclusion.

For instance, a deep learning model might identify a correlation between two seemingly unrelated variables in a dataset. While the model can point out this relationship, it cannot explain why this relationship exists. This lack of interpretability can be problematic, particularly in social sciences, where understanding the 'why' behind patterns is as important as detecting the patterns themselves.

Another challenge is the validation of patterns detected by AI. In traditional social science research, the validation of findings is done through methods such as replication studies or peer review. However, the complexity and non-deterministic nature of AI models can make these methods less feasible. It can be difficult, if not impossible, to precisely replicate an AI model's training due to factors like the randomness in initial weights or the order of training examples. This raises questions about how to validate and trust the patterns that AI uncovers.

Data privacy is another significant challenge in harnessing AI for pattern detection. AI's ability to analyze vast amounts of data from various sources, while a strength, also raises concerns about individuals' privacy rights. There are ethical considerations about what data should be used, how it should be anonymized, and what consent is needed from the individuals whose data is being analyzed.

Furthermore, the use of AI-detected patterns in decision-making processes brings forth concerns of bias and fairness. AI models, despite their sophistication, are susceptible to the adage of "garbage in, garbage out." If the data used to train the model is biased, the patterns the model detects and the decisions it makes will likely be biased as well. This can perpetuate and even exacerbate existing inequalities.

Lastly, there's the challenge of effectively integrating AI within the existing social science research paradigms. There's a need to train social scientists in AI techniques and to create multidisciplinary teams that can leverage both the nuanced human understanding of social scientists and the computational power of AI. There's also a need for regulatory frameworks that guide the use of AI in social sciences, balancing its potential benefits with ethical considerations.

Despite these challenges, the potential of AI to uncover hidden patterns in human behaviour and revolutionize our understanding of social phenomena is enormous. By acknowledging and addressing these challenges, we can responsibly harness this potential and guide the evolution of social sciences in the era of AI.

Ethical Considerations

As we embrace the potential of AI in detecting hidden patterns in human behaviour, it's crucial to carefully navigate the accompanying ethical minefield. Two major areas of concern arise: issues around data privacy and security, and the potential misuse of AI-detected patterns.

Data Privacy and Security

AI systems require vast amounts of data to train and improve. This data often contains sensitive and personal information. The use of such data brings to the forefront concerns about privacy. Do individuals know their data is being used? Have they consented to its use? Is their data being anonymized and securely stored to prevent misuse?

For example, in the case of AI analyzing social media data for patterns in consumer behaviour or mental health, individuals might not be aware that their online activities are being analyzed in this way. Even if they've agreed to the terms of service, the question arises whether they've given informed consent, considering the often-complex and lengthy nature of these agreements.

Data security is another critical concern. With cyber threats becoming increasingly sophisticated, ensuring the security of sensitive data is a constant challenge. Breaches can lead to personal data falling into the wrong hands, with potentially severe consequences for the individuals involved.

Potential for Misuse

Beyond privacy and security, there's a significant concern about how AI-detected patterns could be misused. For instance, patterns in consumer behaviour could be used to manipulate

individuals' buying decisions in ways that may not be in their best interest.

In political scenarios, if AI is used to detect patterns of public sentiment, there's potential for misuse through the spreading of misinformation or propaganda tailored to exploit these patterns.

Furthermore, the misuse of AI-detected patterns can exacerbate social inequalities. If AI models are trained on biased data, they may reinforce these biases in their predictions. For instance, AI used in hiring processes could perpetuate biases against certain demographic groups if it's trained on historical hiring data that reflects these biases.

Lastly, the use of AI-detected patterns in predictive policing or surveillance raises ethical questions. While these applications can aid in crime prevention, they can also infringe on civil liberties and disproportionately target certain communities based on the patterns the AI detects.

In conclusion, while AI holds considerable promise for unearthing hidden patterns in human behaviour, its use needs to be governed by robust ethical frameworks. These should emphasize transparency, consent, fairness, and accountability to ensure that AI serves the best interests of society and doesn't inadvertently amplify inequalities or invade individual privacy.

The Future of AI in Social Sciences

The incorporation of artificial intelligence into social sciences is still in its nascent stages, but the possibilities are exciting and transformative. As AI techniques continue to evolve and mature, we can anticipate some future trends in their application within social sciences.

One key development we can expect to see is the increasing integration of AI and traditional social science methodologies. AI can provide fresh insights into longstanding social phenomena and offer new ways to test and refine theories. Simultaneously, social scientists' expertise will be crucial in guiding AI applications, framing research questions, and interpreting findings.

The trend of "big data" is likely to continue, with increasingly large and diverse data sources being analyzed. Social media, wearables, IoT devices, digital transactions, and more provide a rich tapestry of human behaviour that AI can explore. This will also push for innovations in privacy-preserving technologies, like differential privacy and federated learning, to analyze these datasets responsibly.

We will also likely see advancements in AI interpretability. As researchers strive to resolve the 'black box' problem, new techniques and models will be developed that can provide a more transparent view into AI's decision-making process. This is particularly important in social sciences, where understanding the 'why' behind a pattern or prediction can be as important as the discovery itself.

Additionally, as AI's role in society grows, we can expect more research into the societal impact of AI itself. Social scientists will likely play a key role in studying how AI is influencing our lives, economies, and societies and informing policies to ensure AI's beneficial and equitable use.

However, along with these benefits, there will be potential pitfalls. The misuse of AI and data could lead to privacy infringements, amplification of social biases, or manipulation of consumer or voter behaviour. As AI becomes more ingrained in our societies,

the risk of "AI determinism" – over-reliance on AI at the expense of human judgement and understanding – may also emerge.

It's also worth noting that while AI can help detect patterns, it cannot provide the complete picture. Human behaviour is complex and influenced by myriad factors, many of which may not be captured in the data AI analyzes. Therefore, AI should be seen as a tool that complements, not replaces, traditional social science methodologies.

In conclusion, the future of AI in social sciences promises profound insights and opportunities for societal advancement. At the same time, it calls for ongoing vigilance to navigate ethical, privacy, and interpretability concerns. By balancing these aspects, we can maximize the benefits of AI in understanding and shaping our societies.

Conclusion

As we stand on the precipice of an AI-driven era, understanding the potential of artificial intelligence in uncovering hidden patterns in human behaviour is more important than ever. This exploration has traversed a wide range of aspects, from the fundamental workings of AI to its applications, challenges, ethical considerations, and the potential future landscape in social sciences.

We began by demystifying AI and its inherent capacity to detect complex patterns, showcasing its evolution and increasing sophistication over time. Turning our attention to human behaviour, we acknowledged its multifaceted nature, shaped by diverse influences from cultural to socio-economic factors.

The entry of AI into social sciences has been transformative, as traditional methodologies often struggle to grasp the full

complexity of human behaviour. AI's computational power enables it to unravel non-obvious, intricate patterns that can offer fresh insights into human psychology and sociology. This was vividly demonstrated through various case studies, ranging from consumer behaviour to mental health monitoring.

However, the use of AI in social sciences isn't without challenges. The 'black box' problem, data privacy issues, validation of AI-detected patterns, and potential misuse all pose significant hurdles. These issues, coupled with ethical considerations around data security and potential misuse of AI-detected patterns, call for robust ethical frameworks and informed discourse to ensure responsible application of AI.

Looking forward, the fusion of AI and social sciences promises a future ripe with insight and discovery. However, it is incumbent upon us to ensure the equitable use of AI, respecting privacy, promoting transparency, and preventing undue biases in its applications.

In conclusion, the exploration of hidden patterns in human behaviour using AI holds immense potential to deepen our understanding of ourselves and our societies. As we continue to innovate and develop these technologies, it is crucial to do so responsibly, with an unwavering commitment to advancing human wellbeing and societal progress. The fusion of AI and social sciences not only heralds a new era of discovery but also challenges us to navigate this uncharted territory with wisdom and foresight.

References

1. Russell, S., & Norvig, P. (2016). *Artificial Intelligence: A Modern Approach*. Malaysia: Pearson Education Limited.
2. Dhar, V. (2013). "Data science and prediction". *Communications of the ACM, 56*(12), 64-73.
3. Bzdok, D., & Ioannidis, J. P. (2019). "Exploration, Inference, and Prediction in Neuroscience and Biomedicine". *Trends in neurosciences, 42*(4), 251-262.
4. Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J. F., Breazeal, C., ... & Jennings, N. R. (2019). "Machine behaviour". *Nature, 568*(7753), 477-486.
5. Elish, M. C. (2019). "Moral crumple zones: Cautionary tales in human-robot interaction". *Engaging Science, Technology, and Society, 5*, 40-60.
6. O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books.
7. Miller, T. (2019). "Explanation in artificial intelligence: Insights from the social sciences". *Artificial Intelligence, 267*, 1-38.
8. Castelvecchi, D. (2016). "Can we open the black box of AI?". *Nature News, 538*(7623), 20-23.
9. Crawford, K., & Calo, R. (2016). "There is a blind spot in AI research". *Nature, 538*(7625), 311-313.